

Optimal Time-Convex Hull under the L_p Metrics

Bang-Sin Dai

Dep. of Computer Science
and Information Engineering
National Taiwan Univ.
Taiwan
f94922074@ntu.edu.tw

Mong-Jen Kao

Research Center for
IT Innovation
Academia Sinica
Taiwan
mong@citi.sinica.edu.tw

D.-T. Lee

Dep. of Computer Science
and Engineering
National Chung-Hsing Univ.
Taiwan
dtlee@ieee.org

Abstract

We consider the problem of computing the time-convex hull of a point set under the general L_p metric in the presence of a straight-line highway in the plane. The traveling speed along the highway is assumed to be faster than that off the highway, and the shortest time-path between a distant pair may involve traveling along the highway. The time-convex hull $\text{TCH}(P)$ of a point set P is the smallest set containing both P and *all* shortest time-paths between any two points in $\text{TCH}(P)$. In this paper we give an algorithm that computes the time-convex hull under the L_p metric in optimal $\mathcal{O}(n \log n)$ time for a given set of n points and a real number p with $1 \leq p \leq \infty$.

1 Introduction

Path planning, in particular, shortest time-path planning, in complex transportation networks has become an important yet challenging issue in recent years. With the usage of heterogeneous moving speeds provided by different means of transportation, the *time-distance* between two points, i.e., the amount of time it takes to go from one point to the other, is often more important than their straight-line distance. With the reinterpretation of distances by the time-based concept, fundamental geometric problems such as convex hull, Voronoi diagrams, facility location, etc. have been reconsidered recently in depth and with insights [1, 4, 6].

From the theoretical point of view, straight-line highways which provide faster moving speed and which we can enter and exit at any point is one of the simplest transportation models to explore. The speed at which one can move along the highway is assumed to be $v > 1$, while the speed off the highway is 1. Generalization of convex hulls in the presence of highways was introduced by Hurtado et al. [8], who suggested that the notion of convexity be defined by the inclusion of shortest time paths, instead of straight-line segments, i.e., a set S is said to be **convex** if it contains the shortest time-path between any two points of S . Using this new definition, the time-convex hull $\text{TCH}(P)$ for a set P is the closure of P with respect to the inclusion of shortest time-paths.

In following work, Palop [10] studied the structure of $\text{TCH}(P)$ in the presence of a highway and showed that it is composed of convex clusters possibly together with segments of the highway connecting all the clusters. A particularly interesting fact implied by the hull-structure is that, the shortest time-path between each pair of inter-cluster points must contain a piece of traversal along the highway, while similar assertions do not hold for intra-cluster pairs of points: A distant pair of points (p, q) whose shortest time-path contains a segment of the highway could still belong to the same cluster, for there may exist other points from the same cluster whose shortest time-path to either p or q does not use the highway at all. This suggests that, the structure of $\text{TCH}(P)$ in some sense indicates the degree of convenience provided by

the underlying transportation network. We are content with clusters of higher densities, i.e., any cluster with a large ratio between the number of points of P it contains and the area of that cluster. For sparse clusters, we may want to break them and benefit distant pairs they contain by enhancing the transportation infrastructure.

The approach suggested by Palop [10] for the presence of a highway involves enumeration of shortest time-paths between all pairs of points and hence requires $\Theta(n^2)$ time, where n is the number of points. This problem was later studied by Yu and Lee [11], who proposed an approach based on incremental point insertions in a highway-parallel monotonic order. However, the proposed algorithm does not return the correct hull in all circumstances as particular cases were overlooked. The first sub-quadratic algorithm was given by Aloupis et al. [3], who proposed an $\mathcal{O}(n \log^2 n)$ algorithm for the L_2 metric and an $\mathcal{O}(n \log n)$ algorithm for the L_1 metric, following the incremental approach suggested by [11] with careful case analysis. To the best of our knowledge, no previous results regarding metrics other than L_1 and L_2 were presented.

Our Focus and Contribution. In this paper we address the problem of computing the time-convex hull of a point set in the presence of a straight-line highway under the L_p metric for a given real number p with $1 \leq p \leq \infty$. First, we adopt the concept of *wavefront propagation*, a notion commonly used for path planning [2,6], and derive basic properties required for depicting the hull structure under the general L_p metric. When the shortest path between two points is not uniquely defined, e.g., in L_1 and L_∞ metrics, we propose a re-evaluation on the existing definition of convexity. Previous works concerning convex hulls under metrics other than L_2 , e.g., Ottmann et al. [9] and Aloupis et al. [3], assume a particular path to be taken when multiple choices are available. However, this assumption allows the boundary of a convex set to contain reflex angles, which in some sense deviates from the intuition of a set being convex.

In this work we adopt the definition that requires a convex set to include *every shortest path* between any two points it contains. Although this definition fundamentally simplifies the shapes of convex sets for L_1 and L_∞ metrics, we show that the nature of the problem is not altered when time-based concepts are considered. In particular, the problem of deciding whether any pair of the given points belong to the same cluster under the L_p metric requires $\Omega(n \log n)$ time under the algebraic computation model [5], for all $1 \leq p \leq \infty$.

Second, we provide an optimal $\mathcal{O}(n \log n)$ algorithm for computing the time-convex hull for a given set of points. The known algorithm due to Aloupis et al. [3] stems from a scenario in the cluster-merging step where we have to check for the existence of intersections between a line segment and a set of convex curves composed of parabolae and line segments, which leads to their $\mathcal{O}(n \log^2 n)$ algorithm. In our paper, we tackle this situation by making an observation on the duality of cluster-merging conditions and reduce the problem to the geometric query of deciding if any of the given points lies above a line segment of an arbitrary slope. This approach greatly simplifies the algorithm structure and can be easily generalized to other L_p -metrics for $1 \leq p \leq \infty$. For this particular geometric problem, we use a data structure due to Guibas et al [7] to answer this query in logarithmic time. All together this yields our $\mathcal{O}(n \log n)$ algorithm. We remark that, although our adopted definition of convexity simplifies the shape of convex sets under the L_1 and the L_∞ metrics, the algorithm we propose does not take advantage of this specific property and also works for the original notion for which only a particular path is to be included.

2 Preliminaries

In this section, we give precise definitions of the notions as well as sketches of previously known properties that are essential to present our work. We begin with the general L_p distance metric and basic time-based concepts.

Definition 1 (Distance in the L_p -metrics). For any real number $p \geq 1$ and any two points $q_i, q_j \in \mathbb{R}^n$ with coordinates (i_1, i_2, \dots, i_n) and (j_1, j_2, \dots, j_n) , the distance between q_i and q_j under the L_p -metric is defined to be $d_p(q_i, q_j) = (\sum_{k=1}^n |i_k - j_k|^p)^{\frac{1}{p}}$.

Note that when p tends to infinity, $d_p(q_i, q_j)$ converges to $\max_{1 \leq k \leq n} |i_k - j_k|$. This gives the definition of the distance function in the L_∞ -metric, which is $d_\infty(q_i, q_j) = \max_{1 \leq k \leq n} |i_k - j_k|$. For the rest of this paper, we use the subscript p to indicate the specific L_p -metric, and the subscript p will be omitted when there is no ambiguity.

A *transportation highway* \mathcal{H} in \mathbb{R}^n is a hyperplane in which the moving speed in \mathcal{H} is $v_{\mathcal{H}}$, where $1 < v_{\mathcal{H}} \leq \infty$, while the moving speed off \mathcal{H} is assumed to be *unit*. Given the moving speed in the space, we can define the *time-distance* between any two points in \mathbb{R}^n .

Definition 2. For any $q_i, q_j \in \mathbb{R}^n$, a continuous curve \mathcal{C} connecting q_i and q_j is said to be a *shortest time-path* if the traveling time required along \mathcal{C} is minimum among all possible curves connecting q_i and q_j . The traveling time required along \mathcal{C} is referred to as the time-distance between q_i and q_j , denoted $\hat{d}(q_i, q_j)$.

For any two points q_i and q_j , let $\text{STP}(q_i, q_j)$ denote the set of shortest time-paths between q_i and q_j . For any $\mathcal{C} \in \text{STP}(q_i, q_j)$, we say that \mathcal{C} *enters the highway* \mathcal{H} if $\mathcal{C} \cap \mathcal{H} \neq \emptyset$. The *walking-region* of a point $q \in \mathbb{R}^n$, denoted $\text{WR}(q)$, is defined to be the set of points whose set of shortest time-paths to q contains a time-path that does not enter the highway \mathcal{H} . For any $\mathcal{C} \in \text{STP}(q_i, q_j)$, we say that \mathcal{C} *uses the highway* \mathcal{H} if $\mathcal{C} \cap \mathcal{H}$ contains a piece with non-zero length, i.e., at some point \mathcal{C} enters the highway \mathcal{H} and walks along it.

Convexity and time-convex hulls. In classical definitions, a set of points is said to be *convex* if it contains every line segment joining each pair of points in the set, and the convex hull of a set of points $Q \subseteq \mathbb{R}^n$ is the minimal convex set containing Q . When time-distance is considered, the concept of convexity as well as convex hulls with respect to time-paths is defined analogously. A set of points is said to be *convex with respect to time*, or, *time-convex*, if it contains every shortest time-path joining each pair of points in the set.

Definition 3 (Time-convex hull). The time-convex hull, of a set of points $Q \subseteq \mathbb{R}^n$, denoted $\text{TCH}(Q)$, is the minimal time-convex set containing Q .

Although the aforementioned concepts are defined in \mathbb{R}^n space, in this paper we work in \mathbb{R}^2 plane with an axis-parallel highway placed on the x -axis as higher dimensional space does not give further insights: When considering the shortest time-paths between two points in higher dimensional space, it suffices to consider the specific plane that is orthogonal to \mathcal{H} and that contains the two points.

Time-convex hull under the L_1 and the L_2 metrics. The structure of time-convex hulls under the L_1 and the L_2 metrics has been studied in a series of work [3,10,11]. Below we review important properties. See also Fig. 1 for an illustration.




Figure 1: (a) The only two possible paths for being a shortest time-path between q_1 and q_2 in L_2 . (b)(c) The walking regions of a point $q \in \mathbb{R}^2$ under L_2 and L_1 , respectively.

Proposition 1 ([10, 11]). For the L_2 -metric and any point $q = (x_q, y_q)$ with $y_q \geq 0$, we have the following properties.

1. If a shortest time-path starting from q uses the highway \mathcal{H} , then it must enter the highway with an incidence angle $\alpha = \arcsin 1/v_{\mathcal{H}}$ toward the direction of the destination.
2. The walking region of a point $q \in \mathbb{R}^2$ is characterized by the following two parabola: (a) *right discriminating parabola*, which is the curve satisfying

$$\begin{cases} x \geq x_q + y_q \tan \alpha, & \text{and} \\ \sqrt{(x - x_q)^2 + (y - y_q)^2} = y_q \sec \alpha + y \sec \alpha + \frac{1}{v_{\mathcal{H}}}((x - y \tan \alpha) - (x_q + y_q \tan \alpha)). \end{cases}$$

- (b) The *left discriminating parabola* is symmetric to the right discriminating parabola with respect to the line $x = x_q$.

Proposition 2 ([3]). For the L_1 -metric, the walking region of a point $q = (x_q, y_q)$ with $y_q \geq 0$ is formed by the intersection of the following regions: (a) the vertical strip $x_q - y_q/\beta \leq x \leq x_q + y_q/\beta$, and (b) $y \geq \pm\beta(x - x_q)$, where $\beta = \frac{1}{2}\left(1 - \frac{1}{v_{\mathcal{H}}}\right)$.

3 Hull-Structure under the General L_p -Metrics

In this section, we derive necessary properties to describe the structure of time-convex hulls under the general L_p metrics. First, we adopt the notion of *wavefront propagation* [2, 6], which is a well-established model used in path planning, and derive the behavior of a shortest time-path between any two points. Then we show how the corresponding walking regions are formed, followed by a description of the desired structural properties.

Wavefronts and Shortest Time-Paths. For any $q \in \mathbb{R}^2$, $t \geq 0$, and $p \geq 1$, the wavefront with *source* q and *radius* t under the L_p -metric is defined as

$$W_p(q, t) = \left\{ s : s \in \mathbb{R}^2, \hat{d}_p(q, s) = t \right\}.$$

Literally, $W_p(q, t)$ is the set of points whose time-distances to q are exactly t . Fig. 2 (a) shows the wavefronts, i.e., the “unit-circles” under the L_p metric, or, the p -circles, for different p with $0 < p \leq \infty$ when the highway is not used. The shortest time-path between q and any point $q' \in W_p(q, t)$ is the trace on which q' moves as t changes smoothly to zero, which is a straight-line joining q and q' .

When the highway \mathcal{H} is present and the time-distance changes, deriving the behavior of a shortest time-path that uses \mathcal{H} becomes tricky. Let $q_1, q_2 \in \mathbb{R}^2$, $q_1 \neq q_2$, be two points in the plane, and let $\hat{t}_{1/2}(q_1, q_2) \geq 0$ be the smallest real number such that

$$\text{Bisect}(q_1, q_2) \equiv W_p(q_1, \hat{t}_{1/2}(q_1, q_2)) \cap W_p(q_2, \hat{t}_{1/2}(q_1, q_2)) \neq \emptyset.$$

In other words, $\text{Bisect}(q_1, q_2)$ is the set of points at which $W_p(q_1, t)$ and $W_p(q_2, t)$ meet for the first time. The following lemma shows that $\text{Bisect}(q_1, q_2)$ characterizes the set of “*middle points*” of all shortest time-paths between q_1 and q_2 .

Lemma 1. For each $\mathcal{C} \in \text{STP}(q_1, q_2)$, we have $\mathcal{C} \cap \text{Bisect}(q_1, q_2) \neq \emptyset$. Moreover, for each $q \in \text{Bisect}(q_1, q_2)$, there exists $\mathcal{C}' \in \text{STP}(q_1, q_2)$ such that $q \in \mathcal{C}'$.

Given the set $\text{Bisect}(q_1, q_2)$, a shortest time-path between q_1 and q_2 can be obtained by joining \mathcal{C}_1 and \mathcal{C}_2 , where $\mathcal{C}_1 \in \text{STP}(q_1, q)$ and $\mathcal{C}_2 \in \text{STP}(q, q_2)$ for some $q \in \text{Bisect}(q_1, q_2)$. By expanding the process in a recursive manner we get a set of middle points. Although the cardinality of the set we identified is countable while any continuous curve in the plane contains




Figure 2: (a) p -circles for different values of p : bold rhombus for $p = 1$, bold circle for $p = 2$, and bold square for $p = \infty$. (b) $W_p(q, t)$ for a point $q \in \mathcal{H}$, $v_{\mathcal{H}} < \infty$, and $p > 1$, where the angle α satisfies $\sin \alpha = \frac{\hat{x}_t}{\sqrt{\hat{x}_t^2 + y_p(\hat{x}_t)^2}}$. (c) $W_p(q, t)$ for $v_{\mathcal{H}} < \infty$ and $p = 1$.

uncountably infinite points, it is not difficult to see that, the set of points we locate is *dense*¹ in the underlying curve, and therefore can serve as a representative.

To describe the shape of a wavefront when the highway may be used, we need the following lemma regarding the propagation of wavefronts.

Lemma 2. Let $\mathcal{C}_p^\circ(q, t)$ denote the p -circle with center q and radius t . Then $W_p(q, t)$ is formed by the boundary of

$$\mathcal{C}_p^\circ(q, t) \cup \bigcup_{s: s \in \mathbb{R}^2, \hat{d}_p(q, s) < d_p(q, s) \leq t} W_p(s, t - \hat{d}_p(q, s)).$$

In the following we discuss the case when $1 < p < \infty$ and leave the discussion of shortest time-paths in L_∞ to the appendix for further reference. Let \mathcal{H} be the highway placed on the x -axis with moving speed $v_{\mathcal{H}} > 1$. For any $t \geq 0$, any $0 \leq x \leq t$, and any $1 < p < \infty$, we use $y_p(x, t) = (t^p - |x|^p)^{1/p}$ to denote the y -coordinate of the specific point on the p -circle with x -coordinate x . Let $\hat{x}_t = t \cdot v_{\mathcal{H}}^{1/(1-p)}$. We have the following lemma regarding $W_p(q, t)$. Also refer to Fig. 2 (b) for an illustration.

Lemma 3. For $1 < p < \infty$, $v_{\mathcal{H}} < \infty$, and a point $q \in \mathcal{H}$ which we assume to be $(0, 0)$ for the ease of presentation, the upper-part of $W_p(q, t)$ that lies above \mathcal{H} consists of the following three pieces:

- $\text{Fan}_p(q, t)$: the circular-sector of the p -circle with radius t , ranging from $(-\hat{x}_t, y_p(-\hat{x}_t, t))$ to $(\hat{x}_t, y_p(\hat{x}_t, t))$.
- $\text{LT}_p(q, t)$, $\text{RT}_p(q, t)$: two line segments joining $(-v_{\mathcal{H}} \cdot t, 0)$, $(-\hat{x}_t, y_p(-\hat{x}_t, t))$, and $(\hat{x}_t, y_p(\hat{x}_t, t))$, $(v_{\mathcal{H}} \cdot t, 0)$, respectively. Moreover, $\text{LT}_p(q, t)$ and $\text{RT}_p(q, t)$ are tangent to $\text{Fan}_p(q, t)$.

The lower-part that lies below \mathcal{H} follows symmetrically. For $v_{\mathcal{H}} = \infty$, the upper-part of $W_p(q, t)$ consists of a horizontal line $y = t$.

For each $1 \leq p < \infty$ and $1 < v_{\mathcal{H}} \leq \infty$, we define the real number $\alpha(p, v_{\mathcal{H}})$ as follows. If $p = 1$ or $v_{\mathcal{H}} = \infty$, then $\alpha(p, v_{\mathcal{H}})$ is defined to be zero. Otherwise, $\alpha(p, v_{\mathcal{H}})$ is defined to be

$$\arcsin \frac{v_{\mathcal{H}}^{1/(1-p)}}{\sqrt{v_{\mathcal{H}}^{2/(1-p)} + \left(1 - v_{\mathcal{H}}^{p/(1-p)}\right)^{2/p}}}.$$

Note that, when $p = 2$, this is exactly $\arcsin(1/v_{\mathcal{H}})$. For brevity, we simply use α when there is no ambiguity. The behavior of a shortest time-path that takes the advantage of traversal along the highway is characterized by the following lemma.

Lemma 4. For any point $q = (x_q, y_q)$, $1 \leq p < \infty$, and $1 < v_{\mathcal{H}} \leq \infty$, if a shortest time-path starting from q uses the highway \mathcal{H} , then it must enter the highway with an incidence angle α .

¹ *Dense* is a concept used in classical analysis to indicate that any element of one set can be approximated to any degree by elements of a subset being dense within.

Walking Regions. For any point $q = (x_q, y_q) \in \mathbb{R}^2$ with $y_q \geq 0$, let $q_{\mathcal{H}}^+$ and $q_{\mathcal{H}}^-$ be two points located at $(x_q \pm y_q \tan \alpha, 0)$, respectively. By Lemma 4, we know that, $q_{\mathcal{H}}^+$ and $q_{\mathcal{H}}^-$ are exactly the points at which any shortest time-path from q will enter the highway if needed. This gives the walking region for any point. Let $\alpha = \pi/4$ when $p = \infty$. The following lemma is an updated version of Proposition 1 for general p with $1 \leq p \leq \infty$.

Lemma 5. For any p with $1 \leq p \leq \infty$ and any point $q = (x_q, y_q)$ with $y_q \geq 0$, $\text{WR}_p(q)$ is characterized by the following two curves: (a) *right discriminating curve*, which is the curve $q' = (x', y')$ satisfying $x' \geq x_q + y_q \tan \alpha$ and

$$|\overline{qq'}|_p = |\overline{qq_{\mathcal{H}}^+}|_p + |\overline{q'q_{\mathcal{H}}^-}|_p + \frac{1}{v_{\mathcal{H}}} |\overline{q_{\mathcal{H}}^+ q_{\mathcal{H}}^-}|_p.$$

(b) The *left discriminating curve* is symmetric with respect to the line $x = x_q$.

For any point q , let $\text{WR}_\ell(q)$ and $\text{WR}_r(q)$ denote the left- and right- discriminating curves of $\text{WR}(q)$, respectively. We have the following *dominance property* of the walking regions.

Lemma 6. Let $q_1 = (x_1, y_1)$ and $q_2 = (x_2, y_2)$ be two points such that $x_1 \leq x_2$. If $y_1 \geq y_2$, then $\text{WR}_\ell(q_2)$ lies to the right of $\text{WR}_\ell(q_1)$. Similarly, if $y_1 \leq y_2$, then $\text{WR}_r(q_1)$ lies to the left of $\text{WR}_r(q_2)$.




Figure 3: The left side shows the left boundary of the walking region for the edge $e = (q', q)$ on the right side.

Lemma 6 suggests that, to describe the leftmost and the rightmost boundaries of the walking-regions for a set of points, it suffices to consider the *extreme points*. Let $e = \overline{q_1 q_2}$ be a line segment between two points q_1 and q_2 , where q_1 lies to the left of q_2 . If e has non-positive slope, then the left-boundary of the walking region for e is dominated by $\text{WR}_\ell(q_1)$. Otherwise, we have to consider $\bigcup_{q \in e} \text{WR}_\ell(q)$. By parameterizing each point of e , it is not difficult to see that the left-boundary consists of $\text{WR}_\ell(q_1)$, $\text{WR}_\ell(q_2)$, and their common tangent line. See also Fig. 3 for an illustration.

Closure and Time-Convex Hull of a Point Set. By Lemma 1, to obtain the union of possible shortest time-paths, it suffices to consider the set of all possible bisecting sets that arise inside the recursion. We begin with the closure between pairs of points.

Lemma 7. Let $q_1, q_2 \in \mathbb{R}^2$ be two points. When the highway is not used, the set of all shortest time-paths between q_1 and q_2 is:

- The smallest bounding rectangle of $\{q_1, q_2\}$, when $p = 1$.
- The straight line segment $\overline{q_1 q_2}$ joining q_1 and q_2 , when $1 < p < \infty$.
- The smallest bounding parallelogram whose slopes of the four sides are ± 1 , i.e., a rectangle rotated by 45° , that contains q_1 and q_2 .

Lemma 7 suggests that when the highway is not used and when $1 < p < \infty$, the closure, or, convex hull, of a point set \mathbf{S} with respect to the L_p -metric is identical to that in L_2 , while in L_1 and L_∞ the convex hulls are given by the bounding rectangles and bounding square-parallelograms.




Figure 5: (a) The closure of each cluster under the L_∞ metric when the highway \mathcal{H} is not considered. (b) The closure, i.e., the time-convex hull, for the L_∞ metric.




Figure 4: Time-convex hull for a set of points under the L_p metric where $1 < p < \infty$.

When the highway may be used, the structure of the time-convex hull under the general L_p -metric consists of a set of clusters arranged in a way such that the following holds: (1) Any shortest time-path between intra-cluster pair of points must use the highway. (2) If any shortest time-path between two points does not use the highway, then the two points must belong to the same cluster.

Fig. 4 and Fig. 5 illustrate examples of the time-convex hull for the L_p metrics with $1 < p < \infty$ and $p = \infty$, respectively. Note that, the shape of the closure for each cluster does depend on p and $v_{\mathcal{H}}$, as they determine the incidence angle α .

4 Constructing the Time-Convex Hull

In this section, we present our algorithmic results for this problem. First, we show that, although our definition of convexity simplifies the structures of the resulting convex hulls, e.g., in L_1 and L_∞ , the problem of deciding if any given pair of points belongs to the same cluster already requires $\Omega(n \log n)$ time. Then we present our optimal $\mathcal{O}(n \log n)$ algorithm.

4.1 Problem Complexity

We make a reduction from the *minimum gap problem*, which is a classical problem known to have the problem complexity of $\Theta(n \log n)$. Given n real numbers a_1, a_2, \dots, a_n and a target gap $\epsilon > 0$, the minimum gap problem is to decide if there exist some i, j , $1 \leq i, j \leq n$, such that $|a_i - a_j| \leq \epsilon$.

For any $y \geq 0$, consider the point $q(y) = (0, y)$. Let $\mathcal{C}_{q(y)}: (x, f_{q(y)}(x))$ denote the right discriminating curve of $\text{WR}_p(q(y))$. For any $\epsilon > 0$, let $y_0(p, \epsilon)$ denote the specific real number such that $f_{q(y_0(p, \epsilon))}(\epsilon) = y_0(p, \epsilon)$. Our reduction is done as follows. Given a real number $p \geq 1$ and an instance \mathcal{I} of minimum gap, we create a set \mathbf{S} consisting of n points q_1, q_2, \dots, q_n , where $q_i = (a_i, y_0(p, \epsilon))$ for $1 \leq i \leq n$. The following lemma shows the correctness of this reduction and establishes the $\Omega(n \log n)$ lower bound.

Lemma 8. $y_0(p, \epsilon)$ is well-defined for all p with $1 \leq p \leq \infty$ and all $\epsilon > 0$. Furthermore, the answer to the minimum gap problem on \mathcal{I} is “yes” if and only if the number of clusters in the time-convex hull of \mathbf{S} is less than n .

Corollary 9. Given a set of points \mathbf{S} in the plane, a real number p with $1 \leq p \leq \infty$, and a highway \mathcal{H} placed on the x-axis, the problem of deciding if any given pair of points belongs to the same cluster requires $\Omega(n \log n)$ time.

4.2 An Optimal Algorithm

In this section, we present our algorithm for constructing the time-convex hull for a given point set \mathbf{S} under a given metric L_p with $p \geq 1$. The main approach is to insert the points incrementally into the partially-constructed clusters in ascending order of their x -coordinates. In order to prevent a situation that leads to an undesirably complicated query encountered in the previous work by Aloupis et al. [3], we exploit the symmetric property of cluster-merging conditions and reduce the sub-problem to the following geometric query.

Definition 4 (One-Sided Segment Sweeping Query). Given a set of points \mathbf{S} in the plane, for any line segment L of finite slope, the *one-sided segment sweeping query*, denoted $\mathcal{Q}(L)$, asks if $\mathbf{S} \cap L^+$ is empty, where L^+ is the intersection of the half-plane above L and the vertical strip defined by the end-points of L . That is, we ask if there exists any point $p \in \mathbf{S}$ such that p lies above L .

In the following, we first describe the algorithm and our idea in more detail, assuming the one-sided segment sweeping query is available. Then we show how this query can be answered efficiently.

The given set \mathbf{S} of points is partitioned into two subsets, one containing those points lying above \mathcal{H} and the other containing the remaining. We compute the time-convex hull for the two subsets separately, followed by using a linear scan on the clusters created on both sides to obtain the closure for the entire point set. Below we describe how the time-convex hull for each of the two subsets can be computed.

Let q_1, q_2, \dots, q_n be the set of points sorted in ascending order of their x -coordinates with ties broken by their y -coordinates. During the execution of the algorithm, we maintain the set of clusters the algorithm has created so far, which we further denote by $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$. For ease of presentation, we denote the left- and right-boundary of the walking region of C_i by $\text{WR}_\ell(C_i)$ and $\text{WR}_r(C_i)$, respectively. Furthermore, we use $q \in \text{WR}_\ell(C_i)$ or $q \in \text{WR}_r(C_i)$ to indicate that point q lies to the right of $\text{WR}_\ell(C_i)$ or to the left of $\text{WR}_r(C_i)$, respectively.

In iteration i , $1 \leq i \leq n$, the algorithm inserts q_i into \mathcal{C} and checks if a new cluster has to be created or if existing clusters have to be merged. This is done in the following two steps.

- (a) *Point inclusion test.* In this step, we check if there exists any j , $1 \leq j \leq k$, such that $q_i \in \text{WR}_r(C_j)$. If not, then a new cluster C_{k+1} consisting of the point q_i is created and we enter the next iteration. Otherwise, the smallest index j such that $q_i \in \text{WR}_r(C_j)$ is located. The clusters C_j, C_{j+1}, \dots, C_k and the point q_i are merged into one cluster, which will in turn replace C_j, C_{j+1}, \dots, C_k . Let \mathbf{E} be the set of newly created edges on the upper-hull of this cluster whose slopes are positive. Then we proceed to step (b).
- (b) *Edge inclusion test.* Let k be the number of clusters, and x_0 be the x -coordinate of the leftmost point in C_k . Pick an arbitrary edge $e \in \mathbf{E}$, let \hat{e} denote the line segment appeared on $\text{WR}_\ell(e)$ to which e corresponds, and let $\hat{e}(x_0)$ be the intersection of \hat{e} with the half-plane $x \leq x_0$. Then we invoke the one-sided segment sweeping query $\mathcal{Q}(\hat{e}(x_0))$. If no point lies above $\hat{e}(x_0)$, then e is removed from \mathbf{E} . Otherwise, C_k is merged with C_{k-1} . Let e' be the newly created bridge edge between C_k and C_{k-1} . If e' has positive slope, then it is added to the set \mathbf{E} . This procedure is repeated until the set \mathbf{E} becomes empty.

An approach has been proposed to resolve the *point inclusion test* efficiently, e.g., Yu and Lee [11], and Aloupis et al. [3]. Below we state the lemma directly and leave the technical details to the appendix for further reference.

Lemma 10 ([3, 11]). For each iteration, say i , the smallest index j , $1 \leq j \leq k$, such that $q_i \in \text{WR}_r(C_j)$ can be located in amortized constant time.

To see that our algorithm gives the correct clustering, it suffices to argue the following two conditions: (1) Each cluster-merge our algorithm performs is valid. (2) At the end of each iteration, no more clusters have to be merged.

Apparently these conditions hold at the end of the first iteration, when q_1 is processed. For each of the succeeding iterations, say i , if no clusters are merged in step (a), then the conditions hold trivially. Otherwise, the validity of the cluster-merging operations is guaranteed by Lemma 10 and the fact that if any point lies above \hat{e} , then it belongs to the walking region of e , meaning that the last cluster, C_k , has to be merged again. See also Fig. 6 for an illustration.




Figure 6: When two clusters $C_{q'}$ and C_q are merged and new hull edge e is created, it suffices to check the new walking region e corresponds to, i.e., the dark-gray area.

To see that the second condition holds, let $e = (q', q) \in \mathbf{E}$ be a newly created hull edge, and let $C_{q'}$ and C_q be the two corresponding clusters that were merged. By our assumption that the clusters are correctly created before q_i arrives, we know that the walking-regions of $C_{q'}$ and C_q contain only points that do belong to them, i.e., the light-gray area in the left-hand side of Fig. 6 contains only points from $C_{q'}$ or C_q . Therefore, when $C_{q'}$ and C_q are

merged and e is created, it suffices to check for the existence of points other than C_k inside the new walking region e corresponds to, which is exactly the dark-gray area in Fig. 6. Furthermore, by the dominance property stated in Lemma 6, it suffices to check those edges with positive slopes. This shows that at the end of each iteration when \mathbf{E} becomes empty, no more clusters need to be merged. We have the following theorem.

Theorem 11. Provided that the one-sided segment sweeping query can be answered in $Q(n)$ time using $P(n)$ preprocessing time and $S(n)$ storage, the time-convex hull for a given set \mathbf{S} of n points under the given L_p -metric can be computed in $\mathcal{O}(n \log n + nQ(n) + P(n))$ time using $\mathcal{O}(n + S(n))$ space.

4.2.1 Regarding the One-Sided Segment Sweeping Query.

Below we sketch how this query can be answered efficiently in logarithmic time. Let \mathbf{S} be the set of points, L be the line segment of interest, and \mathcal{I}_L be the vertical strip defined by the two end-points of L . We have the following observation, which relates the query $\mathcal{Q}(L)$ to the problem of computing the upper-hull of $\mathbf{S} \cap \mathcal{I}_L$.

Lemma 12. Let \mathcal{I} be an interval, $\mathcal{C}: \mathcal{I} \rightarrow \mathbb{R}$ be a convex function, i.e., we have $\mathcal{C}(\frac{1}{2}(x_1 + x_2)) \geq \frac{1}{2}(\mathcal{C}(x_1) + \mathcal{C}(x_2)) \forall x_1, x_2 \in \mathcal{I}$, that is differentiable almost everywhere, L be a segment with slope θ_L , $-\infty < \theta_L < \infty$, and $q = (x_q, \mathcal{C}(x_q))$ be a point on the curve \mathcal{C} such that

$$\lim_{x \rightarrow x_q^-} \frac{d\mathcal{C}(x)}{dx} \geq \theta_L \geq \lim_{x \rightarrow x_q^+} \frac{d\mathcal{C}(x)}{dx}.$$

If q lies under \overleftrightarrow{L} , then the curve \mathcal{C} never intersects L .

To help compute the upper-hull of $\mathbf{S} \cap \mathcal{I}_L$, we use a data structure due to Guibas et al [7]. For a given simple path \mathcal{P} of n points with an x -sorted ordering of the points, with $\mathcal{O}(n \log \log n)$ preprocessing time and space, the upper-hull of any subpath $p \in \mathcal{P}$ can be assembled efficiently in $\mathcal{O}(\log n)$ time, represented by a balanced search tree that allows binary search on the hull edges. Note that, q_1, q_2, \dots, q_n is exactly a simple path by definition. The subpath to which \mathcal{I}_L corresponds can be located in $\mathcal{O}(\log n)$ time. In $\mathcal{O}(\log n)$ time we can obtain the corresponding upper-hull and test the condition specified in Lemma 12. We conclude with the following lemma.

Lemma 13. The *one-sided segment sweeping query* can be answered in $\mathcal{O}(\log n)$ time, where n is the number of points, using $\mathcal{O}(n \log n)$ preprocessing time and $\mathcal{O}(n \log \log n)$ space.

5 Conclusion

We conclude with a brief discussion as well as an overview on future work. In this paper, we give an optimal algorithm for the time-convex hull in the presence of a straight-line highway under the general L_p -metric where $1 \leq p \leq \infty$. The structural properties we provide involve non-trivial geometric arguments. We believe that our algorithm and the approach we use can serve as a base to the scenarios for which we have a more complicated transportation infrastructure, e.g., modern city-metros represented by line-segments of different moving speeds. Furthermore, we believe that approaches supporting dynamic settings to a certain degree, e.g., point insertions/deletions, or, dynamic speed transitions, are also a nice direction to explore.

References

- [1] M. Abellanas, F. Hurtado, V. Sacristán, C. Icking, L. Ma, R. Klein, E. Langetepe, and B. Palop. Voronoi diagram for services neighboring a highway. *Inf. Process. Lett.*, 86(5):283–288, June 2003.
- [2] Oswin Aichholzer, Franz Aurenhammer, and Belén Palop. Quickest paths, straight skeletons, and the city Voronoi diagram. In *Proceedings of the eighteenth annual symposium on Computational geometry*, SCG '02, pages 151–159, New York, NY, USA, 2002. ACM.
- [3] Greg Aloupis, Jean Cardinal, Sébastien Collette, Ferran Hurtado, Stefan Langerman, Joseph O'Rourke, and Belén Palop. Highway hull revisited. *Comput. Geom. Theory Appl.*, 43(2):115–130, February 2010.
- [4] Sang Won Bae, Jae-Hoon Kim, and Kyung-Yong Chwa. Optimal construction of the city Voronoi diagram. In *Proceedings of the 17th international conference on Algorithms and Computation*, ISAAC'06, pages 183–192, Berlin, Heidelberg, 2006. Springer-Verlag.
- [5] Michael Ben-Or. Lower bounds for algebraic computation trees. In *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, STOC '83, pages 80–86, New York, NY, USA, 1983. ACM.
- [6] Andreas Gamsa, D.T. Lee, Chih-Hung Liu, and Dorothea Wagner. Higher order city Voronoi diagrams. In *Algorithm Theory SWAT 2012*, volume 7357 of *Lecture Notes in Computer Science*, pages 59–70. Springer Berlin Heidelberg, 2012.
- [7] Leonidas Guibas, John Hershberger, and Jack Snoeyink. Compact interval trees: a data structure for convex hulls. In *Proceedings of the first annual ACM-SIAM symposium on Discrete algorithms*, SODA '90, pages 169–178, Philadelphia, PA, USA, 1990. Society for Industrial and Applied Mathematics.
- [8] F. Hurtado, B. Palop, and V. Sacristán. Diagramas de Voronoi con distancias temporales. Actas de los VIII Encuentros de Geometría Computacional, pages 279–288, 1999. in Spanish.
- [9] Thomas Ottmann, Eljas Soisalon-Soininen, and Derick Wood. On the definition and computation of rectilinear convex hulls. *Information Sciences*, 33(3):157 – 171, 1984.
- [10] B. Palop. *Algorithmic Problems on Proximity and Location under Metric Constraints*. Ph.D thesis, Universitat Politècnica de Catalunya, 2003.
- [11] Teng-Kai Yu and D. T. Lee. Time convex hull with a highway. In *Proceedings of the 4th International Symposium on Voronoi Diagrams in Science and Engineering*, ISVD '07, pages 240–250, Washington, DC, USA, 2007. IEEE Computer Society.

A Hull-Structure in General the L_p -Metrics

A.1 Wavefronts and Shortest Time-Paths

Lemma 1. *For each $\mathcal{C} \in \text{STP}(q_1, q_2)$, we have $\mathcal{C} \cap \text{Bisect}(q_1, q_2) \neq \emptyset$. Moreover, for each $q \in \text{Bisect}(q_1, q_2)$, there exists $\mathcal{C}' \in \text{STP}(q_1, q_2)$ such that $q \in \mathcal{C}'$.*

Proof of Lemma 1. For any $\mathcal{C} \in \text{STP}(q_1, q_2)$ and any $q \in \mathcal{C}$, $\hat{d}(q_1, q)$ is a continuous function taking the values from 0 to $\hat{d}(q_1, q_2)$. Therefore there exists a point $q_{1/2}$ such that $\hat{d}(q_1, q_{1/2}) = \frac{1}{2}\hat{d}(q_1, q_2)$. This implies $\hat{d}(q_{1/2}, q_2) \leq \frac{1}{2}\hat{d}(q_1, q_2)$, where the equality must hold since \mathcal{C} is a shortest time-path between q_1 and q_2 . This shows that $q_{1/2} \in W\left(q_1, \frac{1}{2}\hat{d}(q_1, q_2)\right) \cap W\left(q_2, \frac{1}{2}\hat{d}(q_1, q_2)\right)$. Furthermore, we have $W(q_1, t) \cap W(q_2, t) = \emptyset$ for all $t < \frac{1}{2}\hat{d}(q_1, q_2)$, for otherwise it will result in a shorter time-path than \mathcal{C} , which is a contradiction. Therefore $q_{1/2} \in \text{Bisect}(q_1, q_2)$. This also shows that $\hat{t}_{1/2}(q_1, q_2) = \frac{1}{2}\hat{d}(q_1, q_2)$.

To see the second half holds, for each $q \in \text{Bisect}(q_1, q_2)$, we have $\hat{d}(q_1, q) = \hat{d}(q, q_2) = \hat{t}_{1/2}(q_1, q_2) = \frac{1}{2}\hat{d}(q_1, q_2)$. Joining a shortest time-path from q_1 to q and a shortest time-path from q to q_2 , we get a time-path, say, \mathcal{C}' , whose length is $\hat{d}(q_1, q_2)$. Therefore $\mathcal{C}' \in \text{STP}(q_1, q_2)$. \square

Lemma 2. *$W_p(q, t)$ is formed by the boundary of*

$$\mathcal{C}_p^\circ(q, t) \cup \bigcup_{s: s \in \mathbb{R}^2, \hat{d}_p(q, s) < d_p(q, s) \leq t} W_p\left(s, t - \hat{d}_p(q, s)\right),$$

where $\mathcal{C}_p^\circ(q, t)$ is the p -circle with center q and radius t .

Proof of Lemma 2. To see why this lemma holds, it suffices to observe that, for any $s \in \mathbb{R}^2$ such that $\hat{d}_p(q, s) = d_p(q, s)$, the wavefront of $\mathcal{C}_p^\circ(s, t - \hat{d}_p(q, s))$ is dominated by $\mathcal{C}_p^\circ(q, t)$. \square

Lemma 3. *For $1 < p < \infty$, $v_{\mathcal{H}} < \infty$, and a point $q \in \mathcal{H}$ which we assume to be $(0, 0)$ for the ease of presentation, the upper-part of $W_p(q, t)$ that lies above \mathcal{H} consists of the following three pieces:*

- $\text{Fan}_p(q, t)$: the circular-sector of the p -circle with radius t , ranging from $(-\hat{x}_t, y_p(-\hat{x}_t, t))$ to $(\hat{x}_t, y_p(\hat{x}_t, t))$.
- $\text{LT}_p(q, t)$, $\text{RT}_p(q, t)$: two line segments joining $(-v_{\mathcal{H}} \cdot t, 0)$, $(-\hat{x}_t, y_p(-\hat{x}_t, t))$, and $(\hat{x}_t, y_p(\hat{x}_t, t))$, $(v_{\mathcal{H}} \cdot t, 0)$, respectively. Moreover, $\text{LT}_p(q, t)$ and $\text{RT}_p(q, t)$ are tangent to $\text{Fan}_p(q, t)$.

The lower-part that lies below \mathcal{H} follows symmetrically. For $v_{\mathcal{H}} = \infty$, the upper-part of $W_p(q, t)$ consists of a horizontal line $y = t$.

Proof of Lemma 3. By Lemma 2, it suffices to consider the boundary of

$$\bigcup_{s: \mathcal{H}, \hat{d}_p(q, s) \leq t} W_p\left(s, t - \hat{d}_p(q, s)\right).$$

That is, the boundary of the union of the p -circles $\mathcal{C}(\gamma)$ with center γ and radius $t - \frac{\gamma}{v_{\mathcal{H}}}$, where γ is a real number with $-v_{\mathcal{H}} \cdot t \leq \gamma \leq v_{\mathcal{H}} \cdot t$. See also Fig. 7 (a) for an illustration.

First we argue that the upper-right boundary of

$$\bigcup_{0 \leq \gamma \leq v_{\mathcal{H}} \cdot t} \mathcal{C}(\gamma)$$

contains a line segment that is tangent to $\mathcal{C}(0)$. For each θ with $0 \leq \theta \leq \pi/2$ and each $\mathcal{C}(\gamma)$, let $q(\gamma, \theta)$ denote the specific point of $\mathcal{C}(\gamma)$ such that the line between $q(\gamma, \theta)$ and the center of $\mathcal{C}(\gamma)$ forms an angle of θ above the horizontal line. See also Fig. 7 (b).




Figure 7: The propagation of the p -circles along the highway \mathcal{H} , where in this figure $p = 2.5$.

For a fixed θ , the coordinates of $q(\gamma, \theta)$ is given by the following equations.

$$\begin{cases} y_q = (x_q - \gamma) \cdot \tan \theta \\ (x_q - \gamma)^p + y_q^p = \left(t - \frac{\gamma}{v_{\mathcal{H}}}\right)^p \end{cases}$$

Simplifying the equations, we get

$$y_q = \frac{t - \frac{\gamma}{v_{\mathcal{H}}}}{(1 + \tan^p \theta)^{1/p}} \cdot \tan \theta,$$

which implies that y_q is a linear function of γ . This shows that, for fixed θ , the trace of $q(\gamma, \theta)$ as γ moves from 0 to $v_{\mathcal{H}}t$ is exactly a line segment with the right end-point $(v_{\mathcal{H}}t, 0)$. Since this holds for all θ , and, since we have one-to-one correspondence between θ and a point on the p -circle, union of these segments is exactly the set $\bigcup_{0 \leq \gamma \leq v_{\mathcal{H}}t} \mathcal{C}(\gamma)$. Hence the upper-right boundary will be the upper-most line segment, which is exactly the one tangent to $\mathcal{C}(0)$.

Below, we show that the corresponding tangent point is given by $(\hat{x}_t, y_p(\hat{x}_t, t))$. Consider the curve of the p -circle $\mathcal{C}(0)$ in the first quadrant. The equation is given by $y = (t^p - x^p)^{1/p}$. Hence the slope of the line tangent at any specific point is given by

$$\frac{dy}{dx} = -(t^p - x^p)^{\frac{1-p}{p}} x^{p-1}.$$

The point (x, y) of interest has a tangent line that passes through $(v_{\mathcal{H}}, 0)$. Therefore,

$$\frac{y}{x - v_{\mathcal{H}}} = -(t^p - x^p)^{\frac{1-p}{p}} x^{p-1}.$$

Simplifying the equation we obtain the point $(\hat{x}_t, y_p(\hat{x}_t, t))$ as claimed.

For the case $v_{\mathcal{H}} = \infty$, notice that we have

$$\bigcup_{s: \mathcal{H}, \hat{d}_p(q, s) \leq t} W_p(s, t - \hat{d}_p(q, s)) = \bigcup_{s \in \mathcal{H}} W_p(s, t).$$

The upper-part of the boundary is exactly the horizontal line $y = t$. □

A.2 Regarding the incidence angle, Lemma 4.

Lemma 4. *For any point $q = (x_q, y_q)$, $1 \leq p < \infty$, and $1 < v_{\mathcal{H}} \leq \infty$, if a shortest time-path starting from q uses the highway \mathcal{H} , then it must enter the highway with an incidence angle α .*

Before proving Lemma 4, we need more structural properties regarding the propagation of wavefront under the influence of the highway \mathcal{H} . Let $q = (0, y_q)$ be a point. First, we argue

that, for any $t \geq y_q$ and for all $q' \in \mathcal{H}$ such that $d_p(q, q') \leq t$, $\text{Fan}_p(q', t - d_p(q, q'))$ does not belong to the boundary of $W_p(q, t)$. The trace of each point on $\text{Fan}_p(q', t - d_p(q, q'))$ represents a shortest time-path from q' to that point. Therefore by the triangle inequality, we know that $\text{Fan}_p(q', t - d_p(q, q'))$ will be dominated by the trace of the p -circle with center q .

Below, we prove this lemma for a constrained situation where we have a point q_0 on the highway \mathcal{H} and a point q off the highway.

Lemma 14. For $q_0 = (0, 0) \in \mathcal{H}$ and any $q = (x_q, y_q) \in \mathbb{R}^2$ such that $x_q \geq 0$ and $y_q > 0$, the shortest time-path between q_0 and q is (1) the straight line-segment joining q_0 and q , if $x_q \leq y_q \tan \alpha$, or (2) the straight line-segment from q into highway \mathcal{H} with incidence angle α toward the direction of q_0 , followed by a traversal along the highway to q_0 , if $x_q \geq y_q \tan \alpha$.

Proof. Let $q' = (x_{q'}, y_{q'})$ be a point in $\text{Bisect}(q_0, q)$. We consider the following cases.

- (a) If $q' \in \text{Fan}_p(q_0, \hat{t}_{1/2}(q_0, q))$, then we know that the trace of q' belongs to the p -circle $\mathcal{C}^\circ(q, \hat{t}_{1/2}(q_0, q))$. By the symmetry of p -circles, we know that $\mathcal{C}^\circ(q, \hat{t}_{1/2}(q_0, q))$ and $\text{Fan}_p(q_0, \hat{t}_{1/2}(q_0, q))$ share a common tangent line, and the three points q_0, q, q' are co-linear. Therefore from the trace of p -circles we know that the straight line-segment is a shortest time-path between q_0 and q . Moreover, since the angle $\text{Fan}_p(q_0, \hat{t}_{1/2}(q_0, q))$ spans is between $\pm\alpha$, we have $x_q \leq y_q \tan \alpha$.
- (b) If $q' \in \text{RT}_p(q_0, \hat{t}_{1/2}(q_0, q)) \setminus \mathcal{H}$, then q' must come from the p -circle $\mathcal{C}^\circ(q, \hat{t}_{1/2}(q_0, q))$. Moreover, from the perspective of $\text{RT}_p(q_0, \hat{t}_{1/2}(q_0, q)) \setminus \mathcal{H}$, by the proof of Lemma 3, the trace of the point q' comes from another p -circle from a point, say, $q'' \in \mathcal{H}$ such that q'' and q' forms an angle of α to the vertical line. Therefore the points q, q' , and q'' are again co-linear, and the trace of q' from both sides, which are the line segments $\overline{qq''}$ and $\overline{q''q_0}$, is a shortest time-path between q_0 and q .
- (c) Finally, if $q' \in \mathcal{H}$, then we know that q' does not belong to the p -circle $\mathcal{C}_p^\circ(q, \hat{t}_{1/2}(q_0, q))$ as it will touch \mathcal{H} to the right of q' before it touches q' . Therefore, from (a) we know that $x_q > y_q \tan \alpha + x_{q'}$ for otherwise the trace of the p -circle $\mathcal{C}_p^\circ(q, \hat{t}_{1/2}(q_0, q))$ will touch q' directly. Hence, in the situation between $q' \in \mathcal{H}$ and $q \notin \mathcal{H}$, we know that case (a) will never happen, and we either have case (b) or another recursion. After each iteration, the remaining portion of \mathcal{H} between the two points are shrunk by at least y_q . Therefore we will eventually end up with case (b).

This proves the lemma. □

For any point $q = (x_q, y_q) \in \mathbb{R}^2$ with $y_q \geq 0$, let $q_{\mathcal{H}}^+$ and $q_{\mathcal{H}}^-$ be two points with coordinates $(x_q \pm y_q \tan \alpha, 0)$, respectively. Below, we characterize the behavior of the shortest time-paths between any pair of points. Let $q_1 = (x_1, y_1), q_2 = (x_2, y_2) \in \mathbb{R}^2$ be two points such that $x_1 \leq x_2$ and $y_1, y_2 \geq 0$.

Let q be a point in $\text{Bisect}(q_1, q_2)$. If $q \in \mathcal{H}$, then by the lemma above, we know that there exists shortest time-path that uses the highway \mathcal{H} with incidence angle α . Otherwise, if $q \notin \mathcal{H}$, we have two cases to consider. (1) If q belongs to both $\mathcal{C}_p^\circ(q_1, \hat{t}_{1/2}(q_1, q_2))$ and $\mathcal{C}_p^\circ(q_2, \hat{t}_{1/2}(q_1, q_2))$, then the straight line-segment joining q_1 and q_2 is a shortest time-path between q_1 and q_2 . (2) Otherwise, we know that q belongs to a p -circle and a left-tangent LT_p , which implies the existence of a shortest time-path that uses the highway \mathcal{H} with incidence angle α .

We have considered all possible cases and obtained that, whenever a shortest time-path uses the highway \mathcal{H} , it has an incidence angle α . This proves Lemma 4.

A.3 The L_∞ -metric.

Below we discuss the case when $p = \infty$. While the approach and also the arguments to deriving the properties for the L_∞ metric are similar to that we use for $1 < p < \infty$, we state only the key properties without repeating the technical details. First, regarding the propagation of the wavefront from a point of the highway \mathcal{H} , we have the following lemma.

Lemma 15. For $v_{\mathcal{H}} < \infty$, and a point $q \in \mathcal{H}$ which we assume to be $(0,0)$ for the ease of presentation, the upper-part of $W_\infty(q, t)$ that lies above \mathcal{H} consists of the following three pieces:

- $\text{Fan}_\infty(q, t)$: the circular-sector of the ∞ -circle with radius t , which is a horizontal line segment, ranging from $(-t, t)$ to (t, t) .
- $\text{LT}_\infty(q, t)$, $\text{RT}_\infty(q, t)$: two line segments joining $(-v_{\mathcal{H}} \cdot t, 0)$, $(-t, t)$, and (t, t) , $(v_{\mathcal{H}} \cdot t, 0)$, respectively.

The lower-part that lies below \mathcal{H} follows symmetrically. For $v_{\mathcal{H}} = \infty$, the upper-part of $W_\infty(q, t)$ consists of a horizontal line $y = t$.

Lemma 16. When $v_{\mathcal{H}} < \infty$, if a shortest time-path under the L_∞ metric uses the highway \mathcal{H} , then it must enter the highway with an incidence angle $\alpha = \pi/4$. When $v_{\mathcal{H}} = \infty$, a shortest time-path that uses the highway \mathcal{H} may enter the highway with an incidence angle $-\pi/4 \leq \alpha \leq \pi/4$.

Also refer to Fig. 8 for an illustration on the propagation of wavefronts and the behavior of shortest time-paths under the L_∞ metric.




Figure 8: (a) The wavefront $W_\infty(q, t)$ from a point $q \in \mathcal{H}$. (b) The shortest time-path in L_∞ when the highway is used.

A.4 Walking Regions

Lemma 5. For any p with $1 \leq p \leq \infty$ and any point $q = (x_q, y_q)$ with $y_q \geq 0$, $\text{WR}_p(q)$ is characterized by the following two curve: (a) right discriminating curve, which is the curve $q' = (x', y')$ satisfying $x' \geq x_q + y_q \tan \alpha$ and

$$|qq'|_p = |qq_{\mathcal{H}}^+|_p + |q'q_{\mathcal{H}}^-|_p + \frac{1}{v_{\mathcal{H}}} |q_{\mathcal{H}}^+q_{\mathcal{H}}^-|_p. \quad (1)$$

(b) The left discriminating curve is symmetric with respect to the line $x = x_q$.

Proof of Lemma 5. This lemma follows directly from Lemma 4. Deriving the equation of the curve of the walking-region for p with $1 \leq p < \infty$ is straight-forward. Describing the boundary of the walking-region for L_∞ , however, is a bit more tricky. Below we discuss in more detail.

The two lines with slopes ± 1 that pass through q divide the plane into four regions. Let I, II, and III denote the three specific regions that lie to the left of q from top to bottom. Also refer to Fig. 9 (a) for an illustration.

For points $q' = (x', y') \in \text{I}$, we have $d_\infty(q, q') = y' - y_q$, while $d_\infty(q, \mathcal{H}) + d_\infty(q', \mathcal{H}) = y' + y_q$. This implies any shortest time-path between q and q' will never enter \mathcal{H} . Therefore $\text{I} \subseteq \text{WR}_\infty(q)$. Similarly we have $\text{II} \subseteq \text{WR}_\infty(q)$.

Consider a point $q' = (x', y') \in \text{III}$, we have $d_\infty(q, q') = x_q - x'$. Expanding Eq. (1) we get

$$(x_q - x') = y_q + y' + \frac{1}{v_{\mathcal{H}}}((x_q - y_q) - (x' + y')),$$

which simplifies to

$$y' = x_q - y_q - x'. \quad (2)$$

Fig. 9 (b) shows the boundary of $\text{WR}_\infty(q)$ for a point $q \in \mathbb{R}^2$.




Figure 9: (a) The three regions partitioned by the lines of slope ± 1 . (b) The walking region of a point $q \in \mathbb{R}^2$ under L_∞ .

□

Lemma 6. *Let $q_1 = (x_1, y_1)$ and $q_2 = (x_2, y_2)$ be two points such that $x_1 \leq x_2$. If $y_1 \geq y_2$, then $\text{WR}_\ell(q_2)$ lies to the right of $\text{WR}_\ell(q_1)$. Similarly, if $y_1 \leq y_2$, then $\text{WR}_r(q_1)$ lies to the left of $\text{WR}_r(q_2)$.*

Proof of Lemma 6. First we prove this lemma for the $p = 1$ and $p = \infty$. Assume that $y_1 \geq y_2$.

- For $p = 1$, by Proposition 2, it suffices to consider the relative position between the two lines characterizing the discriminating curve for q_1 and q_2 , respectively. First, since $x_2 \geq x_1$, the line specified in Proposition 2 (b) for q_2 lies to the right of that for q_1 . Similarly, since $y_2 \leq y_1$, we have $(x_2 - y_2/\beta) - (x_1 - y_1/\beta) \geq 0$, which implies that the vertical strip for q_2 also lies to the right of that for q_1 . This implies that $\text{WR}_\ell(q_2)$ lies to the right of $\text{WR}_\ell(q_1)$.
- For $p = \infty$, by Eq. 2 from the proof of Lemma 5, we know that the left discriminating curve for a point $q = (x_q, y_q)$ is characterized by the linear equation $y = x_q - y_q - x$. Since we have $x_1 - y_1 \leq x_2 - y_2$ by our assumption, we know that the left discriminating curve for q_1 lies to the left of that for q_2 .

The other direction of this lemma follows from a symmetric argument.

Below we prove for the case $1 < p < \infty$. To this end, observe that, it suffices to show that for two points $q = (0, y_q)$ and $q' = (0, y_{q'})$ such that $y_q \geq y_{q'}$, $\text{WR}_\ell(q')$ lies to the right of $\text{WR}_\ell(q)$ and $\text{WR}_r(q')$ lies to the left of $\text{WR}_r(q)$.

For the case when $v_{\mathcal{H}} = \infty$, we prove that, for any point $q'_w \in \text{WR}_\ell(q')$, except for the origin $(0, 0)$, there always exists a point $q_w \in \text{WR}_\ell(q)$ with the same y -coordinate such that q_w lies to the left of q'_w . By Lemma 5, any point of $\text{WR}_\ell(q')$ is given by $\text{LT}(q'_{\mathcal{H}}, t) \cap \mathcal{C}^\circ(q', y_{q'} + t)$ for some $t \geq 0$, and there is a one-to-one correspondence.

Let $t_c, t_c \geq 0$, be the specific moment such that

$$\text{LT}(q'_{\mathcal{H}}, t_c) \cap \mathcal{C}^\circ(q', y_{q'} + t_c) = q'_w,$$

and let

$$q_w = \text{LT}(q_{\mathcal{H}}^-, t_c) \cap \mathcal{C}^\circ(q, y_q + t_c).$$

We know that q'_w and q_w share the same y -coordinate as they belong to $\text{LT}(q_{\mathcal{H}}'^-, t_c)$ and $\text{LT}(q_{\mathcal{H}}^-, t_c)$, which are both given by the equation $y = t_c$. Since $y_q \geq y_{q'}$, we know that $\mathcal{C}^\circ(q', y_{q'} + t_c)$ is completely contained by $\mathcal{C}^\circ(q, y_q + t_c)$, which implies that q_w lies to the left of q'_w . This proves that $\text{WR}_\ell(q')$ lies to the right of $\text{WR}_\ell(q)$ for the case $v_{\mathcal{H}} = \infty$. See also Fig. 10 for an illustration.




Figure 10: An illustration of the relative positions between q_w and q'_w for $v_{\mathcal{H}} = \infty$.

For the case $1 < v_{\mathcal{H}} < \infty$, we argue that for any point $q'_w \in \text{WR}_\ell(q') \setminus \mathcal{H}$, there always exists a point $q_w \in \text{WR}_\ell(q)$ such that q_w lies strictly to the lower-left of q'_w . By Lemma 5, any point of $\text{WR}_\ell(q') \setminus \mathcal{H}$ is given by $\text{LT}_p(q_{\mathcal{H}}'^-, t) \cap \mathcal{C}^\circ(q', d_p(q', q_{\mathcal{H}}'^-) + t)$ for some $t > 0$, and there is a one-to-one correspondence. Let t_c , where $t_c > 0$, be the specific moment such that

$$\text{LT}(q_{\mathcal{H}}'^-, t_c) \cap \mathcal{C}^\circ(q', d(q', q_{\mathcal{H}}'^-) + t_c) = q'_w,$$

and let

$$q_w = \text{LT}(q_{\mathcal{H}}^-, t_c) \cap \mathcal{C}^\circ(q, d(q, q_{\mathcal{H}}^-) + t_c).$$

Furthermore, since both $q_{\mathcal{H}}^-$ and $q_{\mathcal{H}}'^-$ belong to \mathcal{H} , it is not difficult to see that $\text{W}(q_{\mathcal{H}}^-, t_c)$ is a horizontal translation of $\text{W}(q_{\mathcal{H}}'^-, t_c)$ to the left-hand side by $d(q_{\mathcal{H}}^-, q_{\mathcal{H}}'^-)$.

To help present our proof, let $\overline{\text{LT}}(q_{\mathcal{H}}'^-, t_c)$ and $\overline{\text{LT}}(q_{\mathcal{H}}^-, t_c)$ denote the lower-part of $\text{W}(q_{\mathcal{H}}'^-, t_c)$ and $\text{W}(q_{\mathcal{H}}^-, t_c)$ that correspond to $\text{LT}(q_{\mathcal{H}}'^-, t_c)$ and $\text{LT}(q_{\mathcal{H}}^-, t_c)$. Moreover, let \bar{q}' and \bar{q} denote the right endpoint of $\overline{\text{LT}}(q_{\mathcal{H}}'^-, t_c)$ and $\overline{\text{LT}}(q_{\mathcal{H}}^-, t_c)$, respectively. See also Fig. 11 for an illustration.

From Lemma 5, the definition of α , and the definition of $q_{\mathcal{H}}'^-$, we know that q' , $q_{\mathcal{H}}'^-$, and \bar{q}' are collinear. Therefore, we know that $\overline{\text{LT}}(q_{\mathcal{H}}'^-, t_c)$ is tangent to $\mathcal{C}^\circ(q', d(q', q_{\mathcal{H}}'^-) + t_c)$ at \bar{q}' . Similarly, $\overline{\text{LT}}(q_{\mathcal{H}}^-, t_c)$ is tangent to $\mathcal{C}^\circ(q, d(q, q_{\mathcal{H}}^-) + t_c)$ at \bar{q} .

Now consider the horizontal translation of $\mathcal{C}^\circ(q', d(q', q_{\mathcal{H}}'^-) + t_c)$ to the left-hand side by $d(q_{\mathcal{H}}^-, q_{\mathcal{H}}'^-)$. Let $\mathcal{C}_{q'}^\circ$ be the translated image. Since $\overline{\text{LT}}(q_{\mathcal{H}}'^-, t_c)$ is tangent to $\mathcal{C}^\circ(q', d(q', q_{\mathcal{H}}'^-) + t_c)$ at \bar{q}' , we know that $\overline{\text{LT}}(q_{\mathcal{H}}^-, t_c)$ is also tangent to $\mathcal{C}_{q'}^\circ$ at \bar{q} , which implies that $\mathcal{C}_{q'}^\circ$ is completely contained within $\mathcal{C}^\circ(q, d(q, q_{\mathcal{H}}^-) + t_c)$. This shows that the translated image of q'_w lies strictly to the upper-right of q_w , since they both belong to $\text{LT}(q_{\mathcal{H}}^-, t_c)$. This implies that q_w lies to the lower-left of q'_w . See also Fig. 12 for an illustration.

The other half, that $\text{WR}_r(q')$ lies to the left of $\text{WR}_r(q)$, can be proved by a symmetric argument. This proves the lemma. \square




Figure 11: An illustration of the relative positions between q , $q_{\mathcal{H}}^-$, and \bar{q} .




Figure 12: An illustration on the translated images of $\mathcal{C}^o(q', d(q', q_{\mathcal{H}}'^-) + t_c)$ and $q'_{w,\leftarrow}$, for which the latter is denoted by $q'_{w,\leftarrow}$.

A.5 Closure and Time-Convex Hull of a Point Set.

Lemma 7. *Let $q_1, q_2 \in \mathbb{R}^2$ be two points, when the highway is not used, the set of all shortest time-paths between q_1 and q_2 , i.e., the closure of the set $\mathbf{S} = \{q_1, q_2\}$ is:*

- *The smallest bounding rectangle of \mathbf{S} , when $p = 1$.*
- *The straight line segment $\overline{q_1 q_2}$ joining q_1 and q_2 , when $1 < p < \infty$.*
- *The smallest bounding parallelogram whose slopes of the four sides are ± 1 , i.e., a rectangle rotated by 45° , that contains \mathbf{S} .*

Proof of Lemma 7. Without loss of generality, assume that $q_1 = (x_1, y_1)$ and $q_2 = (x_2, y_2)$ such that $x_1 \leq x_2$ and $y_1 \leq y_2$.

For the case $p = 1$, observe that both the paths $q_1 \rightarrow (x_1, y_2) \rightarrow q_2$ and $q_1 \rightarrow (x_2, y_1) \rightarrow q_2$ are shortest time-paths between q_1 and q_2 since their lengths are equal to $d_1(q_1, q_2)$. Therefore, the closure of $\mathbf{S} = \{q_1, q_2\}$ contains every shortest time-path between (x_1, y) and (x_2, y) for $y_1 \leq y \leq y_2$. This gives the bounding box of \mathbf{S} , for which we denote by $B_1(\mathbf{S})$ for brevity. Furthermore, it is not difficult to see that, for each $q, q' \in B_1(\mathbf{S})$, $\text{Bisect}(q, q') \subseteq B_1(\mathbf{S})$, since the unit-circles in L_1 are diamonds. Therefore the closure of \mathbf{S} is $B_1(\mathbf{S})$.

Since the unit-circles of L_∞ , squares, are exactly the unit-circles of L_1 being rotated by 45° , we obtain the claimed result for $p = \infty$ by rotating the coordinate system by 45° .

For $1 < p < \infty$, first observe that, any tangent line of a p -circle intersects exactly one point. Furthermore, there exists one-to-one correspondence between the slope of a line and a pair of extreme-points from the p -circle whose joined line segment passes through the center.

Therefore, when the two p -circles with center q_1 and q_2 meet for the first time, they adopt a common tangent line. This shows that, q_1 , q_2 , and the point at which the two p -circles meet are co-linear. The same conclusion holds for further bisecting points. Therefore the shortest time-path between q_1 and q_2 is the straight line-segment between them, and it is unique. This shows that the straight line-segment is also the closure of $\mathbf{S} = \{q_1, q_2\}$. \square

B Constructing the Time-Convex Hull

B.1 Problem Complexity

Lemma 8. $y_0(p, \epsilon)$ is well-defined for all p with $1 \leq p \leq \infty$ and all $\epsilon > 0$. Furthermore, the answer to the minimum gap problem on \mathcal{I} is “yes” if and only if the number of clusters in the time-convex hull of \mathbf{S} is less than n .

Proof of Lemma 8. First we argue that $y_0(p, \epsilon)$ is well-defined. Let $v_{\mathcal{H}} = \infty$. Since for all value of p , the right discriminating curve, $f_{q(y)}$, has a one-on-one correspondence with each value of y , $f^{-1}_{q(y)}$ is well-defined and continuous, spanning value from 0 to infinity as y goes from 0 to infinity. Therefore for all possible ϵ , we can always find a $y_0(p, \epsilon)$. The correctness of this reduction follows directly as testing whether or not a point belongs to the walking region of another is equivalent to the testing of the gap-length. \square

B.2 An Optimal Algorithm

Regarding the Point Inclusion Test. An approach has been proposed to resolve this step efficiently, e.g., Yu and Lee [11], and Aloupis et al. [3]. Below we sketch how it efficiently locates the leftmost cluster that has to be merged with the incoming point q_i .

The main idea is to store $\text{WR}_r(C_j)$ for each cluster C_j created so far by a linked-list, followed by using a linear scan on these lists to retrieve the smallest index for each q_i . The correctness of this approach relies primarily on the following facts:

- For each $1 \leq j, j' \leq k$, $j \neq j'$, $\text{WR}_r(C_j)$ intersects $\text{WR}_r(C_{j'})$ at most once.
- If $q_i \in \text{WR}_r(C_j)$, then all the clusters in between, i.e., C_{j+1}, \dots, C_k , have to be merged as well in order to satisfy the closure property of time-convexity.

What really helps is the *dominance property* implied by the above two properties: Let C_j and $C_{j'}$ be two clusters such that (1) C_j lies to the left of $C_{j'}$ and (2) $\text{WR}_r(C_j)$ intersects $\text{WR}_r(C_{j'})$ at a point q . Then the part of $\text{WR}_r(C_{j'})$ to the right of q is not important any more as it lies above $\text{WR}_r(C_j)$ and whenever a point belongs to $\text{WR}_r(C_j)$, $C_{j'}$ has to be merged as well. This allows us to concentrate on the rightmost boundary of the walking-regions and enables a smooth transition when a new point arrives.

Let $\text{WR}_r(\mathcal{C})$ denote the rightmost boundary of the union of all walking-regions of clusters in \mathcal{C} . During the algorithm, we maintain a set of ordered linked-lists where in each linked-list we store a partial boundary of the walking-regions that belonged to $\text{WR}_r(\mathcal{C})$ at a certain moment. When a new point q_i arrives, we walk the linked-list $\text{WR}_r(\mathcal{C})$ to locate the piece that lies above $\text{WR}_r(q_i)$. This specific piece of $\text{WR}_r(\mathcal{C})$, denoted $\hat{\text{WR}}_r(\mathcal{C})$, is then replaced by the partial piece of $\text{WR}_r(q_i)$ that lies below. The replaced piece, $\hat{\text{WR}}_r(\mathcal{C})$, is inserted to the tail of the ordered linked-lists we maintained as a new list. Then we walk the ordered linked-lists, starting from the tail list, to check if q_i belongs to the walking-region of any cluster and to locate the smallest index if it exists. As the new point always comes in x -monotonic order, line segments and parabolae that are expired with respect to q_i will be discarded during the walk. This walk is continued until we reach a list whose boundary lies above q_i , i.e., it intersects with the vertical

ray shooting upward from q_i . We omit the remaining details and refer the reader to the work of [3, 11].

Theorem 11. *Provided that the one-sided segment sweeping query can be answered in $Q(n)$ time using $P(n)$ preprocessing time and $S(n)$ storage, the time-convex hull for a given set \mathbf{S} of n points under the given L_p -metric can be computed in $\mathcal{O}(n \log n + nQ(n) + P(n))$ time using $\mathcal{O}(n + S(n))$ space.*

Proof of Theorem 11. It remains to argue that the time complexity is $\mathcal{O}(n \log n + nQ(n) + P(n))$. Observe that each cluster-merging operation results in at most one newly-created edge. Therefore at most $\mathcal{O}(n)$ edges is identified and inserted into \mathbf{E} . Furthermore, after each segment query invoked by the algorithm, either one edge is removed from \mathbf{E} , or the number of clusters is decreased by one. Therefore at most $\mathcal{O}(n)$ query is performed. This proves this theorem. \square

Lemma 12. *Let \mathcal{I} be an interval, $\mathcal{C}: \mathcal{I} \rightarrow \mathbb{R}$ be a convex function, i.e., $\mathcal{C}(\frac{1}{2}(x_1 + x_2)) \geq \frac{1}{2}(\mathcal{C}(x_1) + \mathcal{C}(x_2))$, $\forall x_1, x_2 \in \mathcal{I}$, that is differentiable almost everywhere, L be a segment with slope θ_L , $-\infty < \theta_L < \infty$, and $q = (x_q, \mathcal{C}(x_q))$ be a point on the curve \mathcal{C} such that*

$$\lim_{x \rightarrow x_q^-} \frac{d\mathcal{C}(x)}{dx} \geq \theta_L \geq \lim_{x \rightarrow x_q^+} \frac{d\mathcal{C}(x)}{dx}.$$

If q lies under \overleftrightarrow{L} , then the curve \mathcal{C} never intersects L .

Proof of Lemma 12. Notice that, at any point x where $\mathcal{C}(x)$ is differentiable, if $\frac{d\mathcal{C}}{dx}(x) > \theta_L$, then $\mathcal{C}(x)$ tends to go above L as x increases. Similarly, if $\frac{d\mathcal{C}}{dx}(x) < \theta_L$, then $\mathcal{C}(x)$ tends to go below L as x increases.

By the convexity of \mathcal{C} , $\forall x_l, x_r \in \mathcal{I}$ such that $x_l < x_r$, we have

$$\lim_{x \rightarrow x_l^+} \frac{d\mathcal{C}}{dx}(x) \geq \lim_{x \rightarrow x_r^-} \frac{d\mathcal{C}}{dx}(x). \quad (3)$$

Suppose that q lies under \overleftrightarrow{L} and the curve \mathcal{C} intersects L at a point $q' = (x', \mathcal{C}(x'))$. If $x_q < x'$, then we know that at some x'' with $x_q < x'' \leq x'$, $\frac{d\mathcal{C}(x)}{dx} > \theta_L \geq \lim_{x \rightarrow x_q^+} \frac{d\mathcal{C}}{dx}(x)$, which is a contradiction to Ineq. (3) and the fact that $x_q < x'$. For the other case where $x' < x_q$, we obtain a similar contradiction as well. \square